UCRL-CR-147711
B519052

Final Report for "AMG/FOSLS for LLNL Applications"

M. Brezina and J. Jones

September 23, 2001

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

Final Report

University of Colorado Denver
Campus Box 170
Denver, CO 80217-3364

Final Report for "AMG/FOSLS for LLNL Applications"

Submitted by: Marian Brezina, University of Colorado, Denver
PI Name Jim E. Jones, University of California
Title AMG/FOSLS for LLNL Applications

FINAL REPORT
For the period ending September 2001

Prepared for:
University of California
Lawrence Livermore National Laboratory
Attn: Linda Bodtker
P.O. Box 808, L-561
Livermore, CA 94551

Under
B519052
University of Colorado, Denver

Date prepared
September 23, 2001

Final Report

23 September, 2001

1 Outline

The work included improving parallel scalability of the multilevel software package ML, and concentrated on improving robustness and convergence properties of the Smoothed Aggregation Multilevel Iterative Solver (SAMISdat(AMG)) for solution of second order problems. Although both scalar and nonscalar problems were targeted, the emphasis was on treatment of nonscalar problems of thin body elasticity such as encountered with ALE3D.

2 ML smoother improvement

As part of the work at LLNL this summer, the MLS type smoother used by SAMISdat(AMG) has been added as an option to Sandia Lab's ML package.

The MLS smoother is based on a Chebyshev polynomial smoother p minimizing the expression $xp(x)^2$ over the interval $(0, \varrho(A))$ and a first degree Chebyshev polynomial in A $p(A)^2$. Such smoothing procedure has previously been studied in the two-level domain decomposition context [4], where it allowed for a two-level substructuring method with optimal multigrid convergence independent of the coarse-grid size.

The smoother p serves a double-duty in the smoothed aggregation methods as it is also useful in constructing the transfer operators.

The advantage of using MLS smoother is that the serial and parallel implementations are mathematically equivalent, regardless of the number of processors. In addition, its parallelization relies on a matrix-vector product, which is typically easier to parallelize than multiplicative schemes such as the Gauss-Seidel method. Optimized kernel functions exist and can be utilized to expedite development.

In the case of Gauss-Seidel method, the parallelism is often aided by taking a compromise approach in which true Gauss-Seidel is used on each processor, while interprocessor relaxation is performed in an additive fashion similar to

# processors	operator complexity	Gauss-Seidel	MLS
1	1.140	167	161
2	1.175	165	160
4	1.197	166	161
8	1.258	180	150
16	1.335	212	146
32	1.414	266	126
64	1.529	936	123

Table 1: Numbers of iterations in dependence on the number of processors used

block Jacobi method. This is true of the current ML implementation of Gauss-Seidel method, which is not mathematically equivalent to the serial Gauss-Seidel method if multiple processors are used.

As a result, the convergence of the multigrid using the MLS smoother scales better with the number of processors than when ML's default symmetric Gauss-Seidel smoother is used.

This is practically demonstrated in Table 1, which corresponds to an ALE3D benchmark problem of three concentric shells discretized using tetrahedra elements. The problem has 11916 degrees of freedom and features numerical difficulties ranging from high aspect ratios to jumps in material coefficients. For the purpose of testing, we required the Euclidean norm of the residual to be reduced by a factor of 10^{11}. In order to compare the convergence properties with respect to the number of processors, we tested the problem on 1, 2, 4, 8, 16, 32 and 64 processors. The results of this comparison are presented in Table 1. For fairness of comparison, we have used two symmetric Gauss-Seidel iterations as the multigrid smoother, which requires roughly the same amount of computational work as when one iteration of MLS smoother is used.

The results in Table 1 show that the number of iterations necessary to achieve the desired residual reduction grows with the number of processors if the default smoother is used. This tendency is not apparent for small number of processors because the coarsening process used in ML produces higher number of aggregates as the number of processors increases. This leads to somewhat larger operator complexities with increasing processor numbers, which typically results in better convergence rates.

In contrast, method using the MLS smoother benefits from this increase in operator complexity in terms of its convergence rate.

3 Feedback Improvement

It is well known that the convergence properties of a multigrid solver critically depend on the coarse grids being able to accurately represent low energy components of the error.

Formally, we need to guarantee that the transfer operator P between any two levels in the multigrid hierarchy satisfies the weak approximation property:

$$\forall u \in \Re^n \exists v \in \Re^m \quad \text{such that} \quad \|u - Pv\| \leq C_a \|u\|_A. \tag{1}$$

where m, n denote the dimensions of coarse and fine-grid problem, respectively.

In addition, the coarse level basis functions must have small energy. Assuming this is satisfied, we have a convergence rate estimate

$$\|e_{i+1}\|_A \leq (1 - \frac{1}{C})\|e_i\|_A,$$

where C depends on C_a, on the smootness of the coarse space functions, and on the constant in the smoothing property.

In order to satisfy (1), the zero-energy components corresponding to the discretized bilinear form must be resolved exactly by the intergrid transfer operators. For certain problems such as linear elasticity, these are either supplied by the discretization package or can be constructed by the user if nodal coordinates of the discretization nodes are available.

The convergence bounds and the performance of the method can be improved if either the energy of the coarse space basis functions can be reduced or if the constant C_a in the weak approximation property can be made smaller. The smoothed aggregation method provides a way of achieving both simultaneously.

The transfer operators are based on a simple tentative prolongator. The task of the tentative prolongator is to assure satisfaction of the weak approximation property with a good constant. A prolongation smoother is applied to the basis functions given by the tentative prolongator to reduce their energy. The convergence rates can then be improved by either performing more smoothing on the tentative coarse basis or by enriching the set of tentative basis functions.

The range of the prolongator can be enriched by adding explicitly computed eigen-vectors to the range of the prolongator. This has been used in [3] and is also the focus of the active AMGe research [2]. In [1], the coarse space enrichment based on the computed eigenvectors has been combined with adaptively applying more smoothing steps in creation of the transfer operators. All methods of this class achieve good convergence rates for a class of difficult problems, but require local finite element matrices to be available to the solver and tend to be quite costly in terms of storage.

In practice, the local finite element matrices may not be available either because the finite element package does not output them (for example, the Finite Element Interface does not pass this information to the user), or because

the problem may have been discretized by other means than the finite element method. In addition there are situations where the discretization package provides no easy way for passing even the zero-energy modes or the coordinates, or where the nature of the problem does not allow easy construction of the basis for the zero-energy modes based on the given information.

Alterative ways to obtain the required information about the low-energy components corresponding to the bilinear form therefore are of great interest. Our current approach does not rely on local finite element matrices. Instead, a feedback mechanism is used in which the method attempts to determine components of the error which are not effectively reduced by the method at runtime. The transfer operators can then be corrected to allow for these components to be represented. When properly executed, the method should be able to *calibrate* itself for the particular problem.

To achieve this, the setup process had to be rewritten to support the addition of the calibrated "kernel" components. The aggregation procedure had to be revised to make sure that the aggregates support the required number of low energy functions.

At present, the calibration procedure expects an initial vector or a set of zero-energy vectors to be given on input. Based on this information, a multigrid hierarchy can be constructed. In order to recover the low-energy components not represented in the range of the prolongator, we use a small number of iterations of the multigrid method based on the current hierarchy to approximately solve the homegeneous problem $Ax_0 = 0$. If the iteration converges with a convergence rate better than a specified tolerance, we stop the calibration process and use the current hierarchy to solve the problem $Ax = b$. If the convergence stalls, however, we will modify the prolongator in a way which adds x_0 into the rage of the prolongator. The calibration procedure is then restarted to evaluate the quality of the resulting multigrid method.

We allow a maximum of N_c multigrid iterations to be performed within each calibration step. The calibration procedure is terminated if the multigrid convergence rate is better than a given tolerance q_c or if a number of calibration steps has reached a given number C_{steps}.

We have conducted a number of computational experiments on 3D elasticity problems to observe the influence of the calibration procedure on recovering the convergence properties obtained using the full set of rigid body modes (RBM) in the definition of the prolongator. The experiments presented in Tables 2, 3, 4 and 5 treat two ALE3D thin-body elasticity problems. For testing, the specified set of low-energy modes consisted of the three piecewise constant vectors corresponding to the displacement fields. Once the calibration procedure was completed, the problem was solved by applying preconditioned conjugate gradient method with SAMISdat(AMG) as a preconditioner. The PCG iteration was terminated once the Euclidean norm of the residual was reduced by a factor of 10^{11}.

Tables 2 and 3 demonstrate the effects of using the calibration approach to

Prolongator range	iterations	cond
6 RBM	96	86.39
3 displacements + 0 calibrated	230	574.30
$q_c = 0.8, \quad N_c = 5$		
3 displacements + 1 calibrated	198	504.72
3 displacements + 2 calibrated	186	460.01
3 displacements + 3 calibrated	176	395.61
3 displacements + 4 calibrated	163	318.64
3 displacements + 5 calibrated	148	290.63
3 displacements + 6 calibrated	126	209.25
$q_c = 0.9, \quad N_c = 20$		
3 displacements + 1 calibrated	195	494.86
3 displacements + 2 calibrated	170	428.00
3 displacements + 3 calibrated	154	404.15
3 displacements + 4 calibrated	139	273.29
3 displacements + 5 calibrated	115	159.54
3 displacements + 6 calibrated	98	132.09
$q_c = 0.99, \quad N_c = 50$		
3 displacements + 1 calibrated	193	481.26
3 displacements + 2 calibrated	166	434.33
3 displacements + 3 calibrated	148	282.69
3 displacements + 4 calibrated	124	181.08
3 displacements + 5 calibrated	99	133.05
3 displacements + 6 calibrated	69	70.67

Table 2: Convergence properties for problem with 11916 degrees of freedom; 2 SGS pre-smoother and 2 SGS post-smoother used.

improve convergence rate for a ALE3D problem used in Table 1. Tables 4 and 5 present results of similar experiments for the same problem discretized on a finer grid with 158031 degrees of freedom. The first two lines in Tables 2, 3, 4 and 5 corespond to a run based on the full set of 6 rigid body modes and on the piecewise constant displacements, respectively. The following lines correspond to the cases where the coarse space was enriched by adding new vectors obtained through calibration to the piecewise constant coarse space.

The results show that the calibration enrichment can be used to recover the convergence rate obtained using the full set of rigid body modes, provided sufficiently many calibration steps are taken. The experiments also indicate that the calibration procedure using the solver based on MLS smoothing (Tables 3, 5) performed better than that based on the symmetric Gauss-Seidel (Tables 2, 4).

Prolongator range	iterations	cond
6 RBM	76	53.73
3 displacements + 0 calibrated	139	203.97
$q_c = 0.8, \quad N_c = 5$		
3 displacements + 1 calibrated	132	177.37
3 displacements + 2 calibrated	125	163.73
3 displacements + 3 calibrated	114	145.60
3 displacements + 4 calibrated	108	138.84
3 displacements + 5 calibrated	100	121.70
3 displacements + 6 calibrated	92	107.51
$q_c = 0.9, \quad N_c = 20$		
3 displacements + 1 calibrated	125	166.69
3 displacements + 2 calibrated	114	131.50
3 displacements + 3 calibrated	103	111.36
3 displacements + 4 calibrated	91	93.24
3 displacements + 5 calibrated	77	70.15
3 displacements + 6 calibrated	64	41.29
$q_c = 0.99, \quad N_c = 50$		
3 displacements + 1 calibrated	126	166.58
3 displacements + 2 calibrated	111	126.89
3 displacements + 3 calibrated	101	107.59
3 displacements + 4 calibrated	86	82.90
3 displacements + 5 calibrated	63	52.08
3 displacements + 6 calibrated	55	45.71

Table 3: Convergence properties for problem with 11916 degrees of freedom; 1 MLS pre-smoother and 1 MLS post-smoother used.

Prolongator range	iterations	cond
6 RBM	92	99.24
3 displacements + 0 calibrated	213	507.76
$q_c = 0.8, \quad N_c = 5$		
3 displacements + 1 calibrated	178	375.46
3 displacements + 2 calibrated	165	342.32
3 displacements + 3 calibrated	152	296.21
3 displacements + 4 calibrated	138	231.41
3 displacements + 5 calibrated	129	204.05
3 displacements + 6 calibrated	121	192.61
$q_c = 0.9, \quad N_c = 20$		
3 displacements + 1 calibrated	173	361.01
3 displacements + 2 calibrated	153	304.22
3 displacements + 3 calibrated	135	247.16
3 displacements + 4 calibrated	117	187.24
3 displacements + 5 calibrated	105	161.20
3 displacements + 6 calibrated	90	98.75
$q_c = 0.99, \quad N_c = 50$		
3 displacements + 1 calibrated	166	344.44
3 displacements + 2 calibrated	142	245.69
3 displacements + 3 calibrated	109	137.37
3 displacements + 4 calibrated	98	111.16
3 displacements + 5 calibrated	82	79.65
3 displacements + 6 calibrated	74	66.10

Table 4: Convergence properties for problem with 158031 degrees of freedom; 2 SGS pre-smoother and 2 SGS post-smoother used.

Prolongator range	iterations	cond
6 RBM	80	72.57
3 displacements + 0 calibrated	165	320.08
$q_c = 0.8, \quad N_c = 5$		
3 displacements + 1 calibrated	156	256.92
3 displacements + 2 calibrated	142	221.72
3 displacements + 3 calibrated	129	181.45
3 displacements + 4 calibrated	118	159.88
3 displacements + 5 calibrated	107	136.18
3 displacements + 6 calibrated	97	114.37
$q_c = 0.9, \quad N_c = 20$		
3 displacements + 1 calibrated	145	243.35
3 displacements + 2 calibrated	126	180.57
3 displacements + 3 calibrated	109	134.28
3 displacements + 4 calibrated	94	96.49
3 displacements + 5 calibrated	77	71.73
3 displacements + 6 calibrated	66	59.07
$q_c = 0.99, \quad N_c = 50$		
3 displacements + 1 calibrated	145	242.27
3 displacements + 2 calibrated	124	181.14
3 displacements + 3 calibrated	101	119.71
3 displacements + 4 calibrated	86	85.23
3 displacements + 5 calibrated	72	63.59
3 displacements + 6 calibrated	57	35.68

Table 5: Convergence properties for problem with 158031 degrees of freedom; 1 MLS pre-smoother and 1 MLS post-smoother used.

4 Conclusion

Both the parallel performance and our calibration experiments suggest that the MLS smoothing is a viable alternative to using the traditional Gauss-Seidel smoother.

For problems where the discretization software does not provide means for constructing the rigid body modes (e.g., current FEI), the calibration procedure may be necessary to recover good convergence rates, especially with increasing size of the problem. Our computational experiments show that the feedback approach to multigrid setup can recover convergence properties obtained with known set of rigid body modes.

The issue which still needs to be addressed is how to reduce the cost of the current setup procedure, where the multigrid coarse matrices have to be constructed every time a new calibrated vector is added. This cost can, of course, be amortized over multiple solutions when solving a problem with multiple right-hand sides. But it would be desirable to reduce the setup cost even with a single right-hand side.

One possiblility currently under consideration is to use only the tentative setup in the calibration procedure. Another possiblility is to attempt to calibrate several vectors at the same time based on the initial V-cycle. These attempts are in development and will be treated elsewhere.

References

[1] Marian Brezina, Caroline I. Heberton, Jan Mandel, and Petr Vaněk, *An iterative method with convergence rate chosen a priori*, UCD/CCM Report 140, Center for Computational Mathematics, University of Colorado at Denver, February 1999, http://www-math.cudenver.edu/ccmreports/rep140.ps.gz.

[2] Tim Chartier, Rob Falgout, Van Emden Henson, Jim Jones, Tom Manteuffel, Steve McCormick, John Ruge, and Panayot S. Vassilevski, *Spectral AMGe (ρAMGe)*, In preparation (2001).

[3] J. Fish and V. Belsky, *Generalized aggregation multilevel solver*, International Journal For Numerical Methods in Engineering **40** (1997), 4341–4361.

[4] Petr Vaněk, Marian Brezina, and Radek Tezaur, *Two-grid method for linear elasticity on unstructured meshes*, SIAM J. Sci. Comp. **21** (1999), no. 3, 900–923.